On Approximating the Sum-Rate for Multiple-Unicasts

Karthikeyan Shanmugam, Megasthenis Asteris and Alexandros G. Dimakis

Department of Electrical and Computer Engineering
University of Texas at Austin
{karthish, megas}@utexas.edu, dimakis@austin.utexas.edu

Abstract—We study upper bounds on the sum-rate of multiple-unicasts. We approximate the Generalized Network Sharing Bound (GNS cut) of the multiple-unicasts network coding problem with k independent sources. Our approximation algorithm runs in polynomial time and yields an upper bound on the joint source entropy rate, which is within an $O(\log^2 k)$ factor from the GNS cut. It further yields a vector-linear network code that achieves joint source entropy rate within an $O(\log^2 k)$ factor from the GNS cut, but *not* with independent sources: the code induces a correlation pattern among the sources.

Our second contribution is establishing a separation result for vector-linear network codes: for any given field $\mathbb{F}$ there exist networks for which the optimum sum-rate supported by vector-linear codes over $\mathbb{F}$ for independent sources can be multiplicatively separated by a factor of $k^{1-\delta}$, for any constant $\delta > 0$, from the optimum joint entropy rate supported by a code that allows correlation between sources. Finally, we establish a similar separation result for the asymmetric optimum vector-linear sum-rates achieved over two distinct fields $\mathbb{F}_p$ and $\mathbb{F}_q$ for independent sources, revealing that the choice of field can heavily impact the performance of a linear network code.

Index Terms—Multiple unicasts, network coding, index coding, GNS-cut, sum-rate.

I. INTRODUCTION

The *multiple-unicasts network coding problem* is one of the fundamental problems in network information theory. In this problem, k source nodes need to communicate independent information to k corresponding destinations through a directed acyclic network. Information is encoded at the sources and flows through links with limited (typically integral) capacity, while intermediate nodes create (possibly non-linear) combinations of the incoming messages. The canonical question is: *what is the set of transmission rates supported by a given network $\mathcal{G}$ with k independent sources?* A related objective is determining *the optimum achievable sum-rate*, *i.e.*, the optimum joint source entropy rate for the k independent sources. The problem has been extensively studied (see, *e.g.*, [1]–[3] and references therein). It is known that *non-linear* codes are required to achieve the capacity [4], but few papers have studied the question of approximating the rate for multiple unicasts (e.g. [5], [6]).

A significant body of work has focused on developing upper bounds on the joint source entropy rate for multiple-unicasts with independent sources. Several of these bounds belong to the class of *edge cut* bounds, in which the sum-rate is

upper bounded by the cumulative capacity of a appropriately selected set of network links. *Cut set* bounds are a prominent representative of this family, but they are outperformed by a newer member of this class: the *GNS (Generalized Network Sharing) cut bound* [7]. There are several other related bounds including the PdE [8], Information dominance [9] and Functional dependence [10]. With few exceptions (GNS cut and Functional dependence bounds are equivalent), it is not known how these bounds compare. However, all these bounds share one thing in common with the GNS cut: *they are hard to compute*. In this paper we shed new light on the GNS cut bound and its approximation power for the multiple-unicasts network coding problem.

Our Contributions:

1) We tensorize the GNS-cut bound as follows: We use an argument that uses strong graph products to obtain a sequence of rate upper bounds that are valid for vector-linear codes – we show that the weakest bound in this sequence is the GNS cut bound.

2) We define a new communication problem that we call the *relaxed-correlated* multiple-unicasts. In this problem independence across sources is relaxed: the code designer is allowed to introduce *any* correlation structure in the sources in order to maximize the joint source entropy rate. GNS cut is an upper bound on the optimum joint source entropy rate for this relaxed-correlated multiple-unicasts problem.

3) We develop a polynomial time algorithm to provably approximate the GNS cut bound from above within an $O(\log^2 k)$ factor, where k is the number of sources in the network. Our algorithm also yields a *vector-linear code* for the relaxed-correlated sources problem achieving joint source entropy rate within an $O(\log^2 k)$ factor from the optimum over all (even non-linear) network codes.

4) One important question is how the finite field used by the vector-linear code influences the sum-rate. We show that the choice of the field matters tremendously. For *any two fields* $\mathbb{F}_p$ and $\mathbb{F}_q$ and for any $\delta > 0$, there exist multiple-unicast networks for large k such that the optimal sum-rates over $\mathbb{F}_p$ and $\mathbb{F}_q$ differ by a factor of $k^{1-\delta}$, for (Theorem 6). Note that a $1/k$-approximation

can be achieved by having a single unicast and ignoring all other sources. Our result shows that this kind of separation can almost be caused by a poor choice of field. This partially negatively answers an open problem stated recently in [6], asking whether vector-linear codes can approximate the network capacity within a logarithmic factor. Our result shows that the answer is negative for the sum-rate over a fixed field. This relies on a similar result for the symmetric-rates ([11], [12]).

We also show that any field can be bad for some network: for *any given field* $\mathbb{F}$, there exists a multiple-unicasts network (with sufficiently large k) for which the optimum vector-linear joint entropy rates for independent and correlated sources are separated by a factor of $k^{1-\delta}$, for any constant $\delta > 0$ (Theorem 6).

Note that our results do not rule out the approximation of the optimum sum-rate for multiple-unicasts by linear codes in general. They do imply, however, that the achievability must use a field that depends on the network. It is possible that the optimal vector-linear code sum-rate (over the best field for that network) is close to the GNS cut bound and the optimum non-linear sum-rate. Note that previous results imply that this is not possible for the symmetric-rate [11].

Our developments rely on connections between the relaxed-correlated multiple-unicasts network problem and the *index coding problem* established in [13]. In the index coding problem, a single broadcasting agent needs to communicate n distinct messages to n receivers (one message per receiver) over a noiseless broadcast channel. A subset of the source messages is available as side-information to each receiver. The objective is to design a broadcast scheme that uses minimum number of transmissions to deliver the n messages. This has been well studied [14]–[17]. Computing the minimum number of necessary transmissions, even in the case of scalar linear coding schemes, is NP-hard and hard to even approximate within a constant factor [14], [18]. The multiple-unicasts network coding problem can be reduced to an index coding problem when it comes to exact solvability of specific rate tuples [19]. However, under the connection of [19] it is not clear if an approximately 'good' symmetric solution for the reduced index coding problem can be converted to a 'good' network coding solution.

II. DEFINITIONS

We begin with a set of formal definitions that are useful for our subsequent developments.

Definition 1. (Directed Index Coding) *Consider a set of n independent messages (symbols) $\mathbf{x}_i \in \mathbb{F}^p$, $i = 1, \ldots, n$, each consisting of $p \in \mathbb{N}_+$ packets (subsymbols) in some alphabet $\mathbb{F}$, and a set of n users $\{1, \ldots, n\}$, such that user i: 1) wants message x_i, and 2) has messages x_j, $j \in S_i \subseteq \{1, \ldots, n\} \setminus \{i\}$ as side-information. A sender wishes to broadcast all n messages to the corresponding users over a noiseless channel. The objective is to design a coding scheme that minimizes the number of transmissions required for all users to decode their respective messages.* $\diamond$

An Index Coding instance is fully characterized by its *side-information graph* G. The side-information graph G is a directed graph on n vertices corresponding to the n users. An edge (i, j) exists in G if and only if $j \in S_i$, *i.e.*, user i *has* message $\mathbf{x}_j$ as side-information.

Let $\mathbf{x} = [\mathbf{x}_1^T \ \mathbf{x}_2^T \cdots \mathbf{x}_n^T]^T$ be the (pn)-dimensional vector formed by stacking the n symbols $\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{F}^p$. The sender transmits one symbol (or equivalently p-subsymbols) per *channel use*. An $(\mathbb{F}, p, n, r)$ *vector-linear index code* for this problem consists of r linear combinations of symbols in $\mathbf{x}$ over a field $\mathbb{F}$ that satisfies the decodability criterion at every user.

The *broadcast rate* $\beta_{\mathrm{VL}}^{\mathbb{F}}(G, \mathcal{C})$ of an $(\mathbb{F}, p, n, r)$-vector-linear index code $\mathcal{C}$ is the ratio r/p; *the number of channel uses required for all users to receive their message.*[1]

Definition 2. (Multiple-Unicasts (MU) Network) *A multiple-unicasts network instance is an acyclic directed network $\mathcal{G}(\mathcal{N}, \mathcal{E})$ on a set $\mathcal{N}$ of nodes, with the following components:*

1) *$\mathcal{E}$ is the set of links (edges) in the network. Links have unit capacity; they carry at most one bit per channel use. We use $c_{(a,b)}$ to denote the total capacity from node a to node b, i.e., the number of links from a to b. Finally, $h(e)$ and $t(e)$ denote the head and tail of an edge $e \in \mathcal{E}$, respectively.*

2) *(Source/Destination nodes) $\mathcal{S} \triangleq \{s_1, s_2, \ldots, s_k\} \subseteq \mathcal{N}$ is a set of k source nodes, and $\mathcal{T} \triangleq \{t_1, t_2, \ldots, t_k\} \subseteq \mathcal{N}$ is a set of k destination nodes.*

3) *(Source links) $\mathcal{E}_i \subset \mathcal{E}$ is a set of $\mathrm{mincut}(s_i, t_i)$ edges with no tail and head $h(e) = s_i$, $\forall e \in \mathcal{E}_i$, $i = 1, \ldots, k$. Here, $\mathrm{mincut}(s_i, t_i)$ is the number of unit-capacity links in the minimum cut between source s_i and destination t_i. We refer to $\mathcal{E}_i$ as the set of source links of source s_i.*

Each source node s_i wants to transmit information to its corresponding destination t_i, $i = 1, \ldots, k$. Information is fed into the network through the source links $\cup_{i=1}^{k} \mathcal{E}_i$. $\diamond$

The *multiple-unicasts network coding problem* is the problem of designing a *network code*: the set of rules that govern how information is encoded and flows through the network. One of the canonical objectives of *multiple-unicasts network coding* is to maximize the total amount of information transmitted through the network per channel use, *i.e.*, to maximize the *joint source entropy rate*. Here, we focus only on *vector-linear codes*, *i.e.*, codes in which encoding and decoding involve only vector-linear operations.

Definition 3. (Vector Linear MU Network Code) *An $(\mathbb{F}, p, m, r)$ vector-linear MU network code $\mathcal{C}$ is a collection of vectors $\mathbf{z}_e \in \mathbb{F}^p, \forall e \in \mathcal{E}$ that depend on the aggregate source message vector $\mathbf{x} \in \mathbb{F}^r$ (consisting of r independent subsymbols) satisfying:*

---

[1] Recall that a channel use is the transmission of a symbol, or equivalently the transmission of p-subsymbols.

1) Coding at intermediate nodes: *For source link e, $\mathbf{z}_e$ is a linear combination of sub-symbols in $\mathbf{x}$. For each non-source link $e \in \mathcal{E}$, $\mathbf{z}_e$ is a linear combination of $\mathbf{z}_a$'s of the edges incident on it, i.e. $\{\mathbf{z}_a\}_{a:h(a)=t(e)}$.*

2) Decoding at destinations: *At every destination t_i, every variable $\mathbf{z}_e$ for $e \in \mathcal{E}_i$, is linearly decodable from information flowing into t_i, i.e. $\{\mathbf{z}_a\}_{a:t(a)=t_i}$.*

3) Independence between sources: *The variables of one source, i.e. $\{\mathbf{z}_e\}_{e \in \mathcal{E}_i}$ are mutually independent of those of other sources.*

The joint source entropy rate achieved by such a code is equal to r/p bits per channel use. Due to the independence among sources, the joint source entropy rate is equal to the *sum-rate* of the k sources. We use $R^{\text{MU}}(\mathcal{G}; \mathbb{F})$ to denote the *optimum sum-rate achievable over all vector-linear network codes defined over the field* $\mathbb{F}$, and $R^{\text{MU}}(\mathcal{G})$ to denote the optimum vector-linear sum-rate over all fields.

Relaxed-Correlated Sources. For our developments it is useful to consider a variant of the multiple-unicasts network coding problem, in which the requirement that source information is independent across sources is overlooked. We refer to the modified version as the problem of *relaxed-correlated sources*. In the modified problem, we still seek to maximize the maximum joint source entropy, but allow arbitrary correlations among sources.

Definition 4. (**Vector-Linear Relaxed-Correlated MU Network Code**) *A vector-linear code $\mathcal{C}$ for the multiple-unicasts network coding problem with* relaxed-correlated sources *is defined as in Def. 3 omitting requirement* (3). ◇

We use $R^{\text{co}}(\mathcal{G}; \mathbb{F}) = r/p$ to denote the optimum joint source entropy rate achievable by vector-linear codes over a given field $\mathbb{F}$ in the *relaxed-correlated sources* problem, and $R^{\text{co}}(\mathcal{G})$ to denote the optimum rate over all fields, accordingly. Clearly, $R^{\text{MU}}(\mathcal{G}) \leq R^{\text{co}}(\mathcal{G})$.

Remark 1. *We emphasize that the optimal joint source entropy rate $R^{\text{co}}(\mathcal{G})$ in the relaxed-correlated sources problem is achieved for some (unspecified) source correlation pattern. The correlation pattern is an additional degree of freedom towards maximizing the joint source entropy rate; it is* not *a code design constraint.*

Remark 2. *In Def. 2, we require $|\mathcal{E}_i| = \text{MINCUT}(s_i, t_i)$. This is only a useful convention and does not affect the value of $R^{\text{MU}}(\mathcal{G})$. It does, however, affect $R^{\text{co}}(\mathcal{G})$. In this work, we upper bound $R^{\text{MU}}(\mathcal{G})$ by developing bounds on $R^{\text{co}}(\mathcal{G})$. Hence, the convention becomes essential.*

III. BOUNDS ON THE SUM-RATE OF AN MU NETWORK

We develop upper bounds on $R^{\text{MU}}(\mathcal{G})$, the optimum sum-rate supported by an MU network with independent sources using vector-linear codes. In fact, our bounds are developed for $R^{\text{co}}(\mathcal{G})$, the optimum vector-linear joint source entropy rate in the relaxed-correlated sources problem. Our results extend those in [13].

A. From Multiple-Unicasts Network Coding to Index Coding

Consider a multiple-unicasts network $\mathcal{G}$ with k sources and m links. Let $\mathcal{G}'$ be a directed cyclic network constructed from $\mathcal{G}$ by setting $t(e) = t_i$, $\forall e \in \mathcal{E}_i$, $i = 1, \ldots, k$, *i.e.*, setting the destination node t_i to be the tail of every source link of source s_i.

Let G be the *(reversed)* [2] *line graph* of $\mathcal{G}'$, *i.e.*, a directed graph on m vertices corresponding to the m links in $\mathcal{G}'$, with a directed edge from vertex v to $\widehat{v}$ corresponding to links e and $\widehat{e}$, respectively, iff $h(e) = t(\widehat{e})$ in $\mathcal{G}'$.

Theorem 1 ([13]). *Consider a multiple-unicasts network $\mathcal{G}$ with m links, and a vector-linear code $\mathcal{C}$ with correlated sources, achieving joint source entropy rate r. The dual code $\mathcal{C}^{\perp}$ is a vector-linear index code achieving rate $m - r$ in the index coding instance with side-information graph G constructed based on $\mathcal{G}$ as described in Section III-A.*

Corollary 1. *If G is the directed graph constructed based on the network $\mathcal{G}$ as described in Section III-A, then*

$$R^{\text{co}}(\mathcal{G}) = m - \beta_{\text{VL}}(G).$$

We exploit the connection established in Cor. 1 to develop upper bounds on the joint source entropy rate $R^{\text{co}}(\mathcal{G})$, through properties of the side information graph G of the associated index coding problem.

Definition 5. $\text{MAIS}(G)$ *of a directed graph G is the cardinality of the largest set $\widehat{V} \subseteq V(G)$ such that the subgraph of G induced by $\widehat{V}$ is acyclic.*

It is known that the size of the maximum acyclic subgraph of G is a lower bound on $\beta_{\text{VL}}(G)$. Tighter bounds can be obtained via graph tensorization.

Lemma 1. *The optimum broadcast rate $\beta_{\text{VL}}(G)$ of an index coding instance with side-information graph G, satisfies*

$$\sqrt[q]{\text{MAIS}(\otimes^q G)} \leq \beta_{\text{VL}}(G), \quad \forall q \in \mathbb{Z}^+,$$

where $\otimes^q$ denotes the strong product of G with itself q times.

Proof: See proof of Theorem 9 in the Appendix. ■

Theorem 2. *Consider a multiple-unicasts network $\mathcal{G}$ with k sources and m links. Further, let G be the digraph on m vertices obtained from $\mathcal{G}$ as described in Section III-A. Then,*

$$R^{\text{MU}}(\mathcal{G}) \leq R^{\text{co}}(\mathcal{G}) = m - \beta_{\text{VL}}(G)$$
$$\leq m - \sqrt[q]{\text{MAIS}(\otimes^q G)}, \quad q \in \mathbb{Z}^+.$$

Proof: The proof follows from Thm. 1 and Lem. 1. ■

B. Comparison with GNS cut bound

We compare the bounds of Theorem 2 with the GNS cut bound; the weakest among the former (*i.e.*, for $q = 1$) is at least as good as the latter. A more careful application of the GNS cut approach (on a slightly modified network) reveals that

---

[2]We refer to G as the *reversed* line graph of $\mathcal{G}$ because the direction of its edges is reversed compared to the typical definition of a line graph.

the two bounds are effectively equal. First, recall the definition of the GNS cut:

Definition 6 ([7]). *A GNS cut of a multiple-unicasts network* $\mathcal{G}(\mathcal{V}, \mathcal{E})$ *with* k *sources, is a subset* $S \subset \mathcal{E}$ *such that for* $\mathcal{G} - S$ *(i.e., the network obtained by removing the links in S from* $\mathcal{G}$*) the following holds: there exists a permutation* $\pi : [k] \to [k]$ *such that* $\forall i, j \in [k]$, *if* $\pi(i) \geq \pi(j)$, *then no path exists from source* s_i *to destination* t_j.

The size of the smallest (in terms of capacity) GNS cut, denoted by $\mathsf{GNSCUT}(\mathcal{G})$, is an upper bound on the *non-linear* sum-rate of the multiple-unicasts problem with independent sources [7].

Theorem 3. *Consider a multiple-unicasts network* $\mathcal{G}$ *with* k *sources and* m *links. Let* G *be a digraph on* m *vertices constructed based on* $\mathcal{G}$ *as described in Section III-A. Then,*

$$m - \mathsf{MAIS}(G) \ \leq \ \mathsf{GNSCUT}(\mathcal{G}).$$

Proof: See Appendix, Section A. ∎

In other words, the bounds of Theorem 2 are at least as tight as the GNS cut bound.

The GNS cut technique can be slightly strengthened to yield an upper bound exactly equal to $m - \mathsf{MAIS}(G)$. We achieve that by obtaining the GNS cut bound on a modified, yet equivalent network. Given a multiple-unicasts network $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with k sources and m links, consider a network $\widetilde{\mathcal{G}}(\widetilde{\mathcal{V}}, \widetilde{\mathcal{E}})$ obtained from $\mathcal{G}$ as follows:

1) Introduce k nodes $\widetilde{s}_1, \ldots, \widetilde{s}_k$ to $\mathcal{G}$, i.e., $\widetilde{\mathcal{V}} = \mathcal{V} \cup \{\widetilde{s}_i\}_{i=1}^{k}$.
2) Set $t(e) = \widetilde{s}_i$, $\forall e \in \mathcal{E}_i$, $i = 1, \ldots, k$, that is, set $\widetilde{s}_i$ as the tail of all source links of source s_i.
3) Introduce a set $\widetilde{\mathcal{E}}_i$ of $|\mathcal{E}_i|$ new links with head $\widetilde{s}_i$ and no tail, for all $i \in \{1, \ldots, k\}$.

The modified network $\widetilde{\mathcal{G}}$ is a multiple-unicasts network with k sources $\widetilde{s}_1, \ldots, \widetilde{s}_k$ and respective destinations $t_1, \ldots, t_k$. One can verify that $R^{\text{co}}(\widetilde{\mathcal{G}}) = R^{\text{co}}(\mathcal{G})$. The key difference is that the $|\mathcal{E}_i|$ source links of source s_i in $\mathcal{G}$ have become regular links in $\widetilde{\mathcal{G}}$ and can be used in a GNS cut. Thus, the bound obtained on the modified network is potentially tighter, *i.e.*, $\mathsf{GNSCUT}(\widetilde{\mathcal{G}}) \leq \mathsf{GNSCUT}(\mathcal{G})$.

Theorem 4. *Consider a multiple-unicasts network* $\mathcal{G}$ *with* k *sources and* m *links. Let* G *be the digraph on* m *vertices obtained from* $\mathcal{G}$ *as described in Section III-A, and* $\widetilde{\mathcal{G}}$ *the modified network constructed as described above. Then, any feasible feedback vertex set of* G *corresponds to a GNS cut in* $\widetilde{\mathcal{G}}$ *with the same capacity. In turn,*

$$m - \mathsf{MAIS}(G) = \mathsf{GNSCUT}(\widetilde{\mathcal{G}}).$$

Proof: The proof relies on showing that each GNS cut in $\widetilde{\mathcal{G}}$ corresponds to a *Feedback Vertex Set* (FVS) of equal size in the digraph G. A FVS $\mathcal{F}$ is a set of vertices such that the subgraph of G induced by $V(G) - \mathcal{F}$ is acyclic. By definition, $m - \mathsf{MAIS}(G)$ is the cardinality of the *minimum feedback vertex set* in G.

See Appendix, Section B for the complete proof. ∎

Remark 3. $\mathsf{GNSCUT}(\mathcal{G})$ *is an upper bound on the* non-linear *sum-rate of the multiple-unicasts network coding problem on* $\mathcal{G}$ *with independent sources (Thm.* 1 *in Chapter* 2 *of [7], Thm.* 2 *in [20]). Those results can be generalized to show that* $\mathsf{GNSCUT}(\mathcal{G})$ *also upper bounds the non-linear joint source entropy rate in the problem of relaxed-correlated sources. By Theorem 4, it follows that* $m - \mathsf{MAIS}(G)$ *is also an upper bound on the optimum* non-linear *joint source entropy rate in both problems.*

IV. APPROXIMATING THE GNS CUT BOUND

Determining the GNS cut bound for a given network is computationally hard problem in general [7]. We describe an algorithm to approximately compute the GNS cut bound for a given acyclic network $\widetilde{\mathcal{G}}$. We exploit the connection of Theorem 4 and the special structure of a multiple-unicasts network $\widetilde{\mathcal{G}}$, and we utilize known approximation algorithms for the *Feedback Vertex Set problem* on a diagraph.

The Feedback Vertex Set (FVS) problem, *i.e.*, the problem of finding the smallest FVS in a given digraph G, is NP-complete [21]. The LP dual of its LP relaxation is the *fractional cycle packing problem* [21], [22]. A fractional cycle packing is a function $q(C)$ from the set of cycles $\mathcal{C}$ in G to $[0, 1]$, satisfying $\sum_{C \in \mathcal{C}: v \cap C \neq \emptyset} q(C) \leq 1$, for each $v \in V(G)$. Letting $|q| = \sum_{C \in \mathcal{C}} q(C)$, the *fractional cycle packing number* $r_{\text{CP}}(G)$ of G is defined to be the maximum of $|q|$ taken over all fractional cycle packings q in G. Clearly, $r_{\text{CP}}(G) \leq |\mathcal{F}|$ for *all* feedback vertex sets $\mathcal{F}$ in G. By definition, $m - \mathsf{MAIS}(G)$ is the cardinality of the *minimum feedback vertex set* in G. Therefore,

$$r_{\text{CP}}(G) \ \leq \ m - \mathsf{MAIS}(G).$$

An *optimal* fractional cycle packing [23] [22] (or an $(1 + \epsilon)$ approximation, $\epsilon > 0$) can be computed in polynomial time (in m, ϵ^{-1}). A feasible fractional cycle packing on G can be suitably rounded to yield a FVS $\mathcal{F}$ with cardinality $|\mathcal{F}| \leq r_{\text{CP}}(G) \cdot O(\log m \log \log m)$. We conclude that for any directed graph G on m vertices,

$$m - \mathsf{MAIS}(G) \ \leq \ r_{\text{CP}}(G) \cdot O(\log m \log \log m). \quad (1)$$

Note that (1) holds for arbitrary digraphs. But G has special structure as it is the (reverse) line-graph of a multiple-unicasts network $\mathcal{G}'$ (itself a modification of a network $\mathcal{G}$) as described in Section III-A. Any feedback *vertex* set in G maps trivially to a feedback *edge* set in $\mathcal{G}'$. All cycles in $\mathcal{G}'$ go through the k source nodes of $\mathcal{G}$. There exist polynomial-time algorithms that exploit this additional structure to compute a feedback edge set (in turn, a feedback vertex set in G) with cardinality within a $O(\log^2 k)$ factor from $r_{\text{CP}}(G)$ [22].

Theorem 5. *Consider a multiple-unicasts network* $\mathcal{G}$ *with* k *sources and* m *unit-capacity links. Let* G *be the digraph on* m *vertices obtained from* $\mathcal{G}$ *as described in Section III-A. Then,*

$$r_{\text{CP}}(G) \ \leq \ m - \beta_{\text{VL}}(G)$$
$$\leq \ m - \mathsf{MAIS}(G) \ \leq \ r_{\text{CP}}(G) \cdot O(\log^2 k),$$

where $r_{\text{CP}}(G)$ is the fractional cycle packing number *of G. Further, $r_{\text{CP}}(G)$ also equals the joint source entropy rate supported by a feasible (and polynomial-time computable) vector-linear multiple-unicasts network code for the relaxed-correlated sources problem on $\mathcal{G}$.*

Proof: See Appendix, Section C. ∎

V. PRICE OF INDEPENDENCE

The GNS cut, similar to the novel bounds of Theorem 2, upper bound the optimum vector-linear joint source entropy rate for the relaxed-correlated sources, and in turn for independent sources since $R^{\text{CO}}(\mathcal{G}) \geq R^{\text{MU}}(\mathcal{G})$. However, it remains unclear how the gap between the two rates scales. The following Theorem takes a step towards addressing this question.

Theorem 6. *For any prime field $\mathbb{F}_p$, for any constant $\delta > 0$, there is a k sufficiently large and there exists a family of multiple-unicasts network instances $\mathcal{G}$ with k sources (k sufficiently large) for which $R^{\text{CO}}(\mathcal{G}; \mathbb{F}_p) \geq k^{1-\delta} \cdot R^{\text{MU}}(\mathcal{G}; \mathbb{F}_p)$. Further, for any two fields $\mathbb{F}_p$ and $\mathbb{F}_q$, for any $\delta > 0$, there is a large enough k and a multiple-unicasts network $\mathcal{G}$ such that $R^{\text{MU}}(\mathcal{G}; \mathbb{F}_q) \geq k^{1-\delta} \cdot R^{\text{MU}}(\mathcal{G}; \mathbb{F}_p)$, .*

Proof: See Appendix, Section D. ∎

Theorem 6 effectively states that for a *fixed* field, there exists networks for which the optimum sum-rate over all vector-linear codes *over that field* is almost a k-factor away from the GNS cut bound. Second, when designing a vector-linear code for a given multiple-unicasts network, the choice of field can have a tremendous impact on performance: it can affect the achievable sum-rate by almost a factor of k.

VI. CONCLUSIONS

We presented a sequence of upper bounds on the sum-rate for multiple-unicasts that are valid for vector-linear codes. The first bound in this sequence is equivalent to the GNS cut bound, which also holds for non-linear codes. Further, we showed that the GNS cut bound can be approximated within an $O(\log^2 k)$ factor in polynomial time. This is, to the best of our knowledge, the only case in the family of generalized cut-set bounds [7] that can be efficiently approximated. Finally, we show the importance of the field used by the vector-linear code: the GNS cut and the capacity can be very far from the best vector-linear code over a poorly chosen field.

REFERENCES

[1] T. Ho and D. Lun, *Network coding: an introduction.* Cambridge University Press, 2008.
[2] A. Ramakrishnan, A. Das, H. Maleki, A. Markopoulou, S. Jafar, and S. Vishwanath, "Network coding for three unicast sessions: Interference alignment approaches," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on.* IEEE, 2010, pp. 1054–1061.
[3] C. Meng, A. Ramakrishnan, A. Markopoulou, and S. A. Jafar, "On the feasibility of precoding-based network alignment for three unicast sessions," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on.* IEEE, 2012, pp. 1907–1911.
[4] R. Dougherty, C. Freiling, and K. Zeger, "Networks, matroids, and non-shannon information inequalities," *Information Theory, IEEE Transactions on*, vol. 53, no. 6, pp. 1949–1969, 2007.
[5] M. Langberg and A. Sprintson, "On the hardness of approximating the network coding capacity," *Information Theory, IEEE Transactions on*, vol. 57, no. 2, pp. 1008–1014, 2011.
[6] S. Lovett, "Linear codes cannot approximate the network capacity within any constant factor." *ECCC TR14-141*, 2014.
[7] S. Kamath, "A study of some problems in network information theory," Ph.D. dissertation, EECS Department, University of California, Berkeley, Aug 2013.
[8] G. Kramer and S. A. Savari, "Edge-cut bounds on network coding rates," *Journal of Network and Systems Management*, vol. 14, no. 1, pp. 49–67, 2006.
[9] N. J. Harvey, R. Kleinberg, and A. Lehman, "On the capacity of information networks," *Information Theory, IEEE Transactions on*, vol. 52, no. 6, pp. 2345–2364, 2006.
[10] S. Thakor, A. Grant, and T. Chan, "Network coding capacity: A functional dependence bound," in *Information Theory, 2009. ISIT 2009. IEEE International Symposium on.* IEEE, 2009, pp. 263–267.
[11] A. Blasiak, "A graph-theoretic approach to network coding," Ph.D. dissertation, Cornell University, 2013.
[12] E. Lubetzky and U. Stav, "Nonlinear index coding outperforming the linear optimum," *Information Theory, IEEE Transactions on*, vol. 55, no. 8, pp. 3544–3551, 2009.
[13] K. Shanmugam and A. G. Dimakis, "Bounding multiple unicasts through index coding and locally repairable codes," *arXiv preprint arXiv:1402.3895*, 2014.
[14] Z. Bar-Yossef, Y. Birk, T. Jayram, and T. Kol, "Index coding with side information," *Information Theory, IEEE Transactions on*, vol. 57, no. 3, pp. 1479–1494, 2011.
[15] Y. Birk and T. Kol, "Informed-source coding-on-demand (iscod) over broadcast channels," in *INFOCOM'98. Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 3. IEEE, 1998, pp. 1257–1264.
[16] H. Maleki, V. Cadambe, and S. Jafar, "Index coding: an interference alignment perspective," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on.* IEEE, 2012, pp. 2236–2240.
[17] A. Blasiak, R. Kleinberg, and E. Lubetzky, "Broadcasting with side information: Bounding and approximating the broadcast rate," *Information Theory, IEEE Transactions on*, vol. 59, no. 9, pp. 5811–5823, 2013.
[18] R. Peeters, "Orthogonal representations over finite fields and the chromatic number of graphs," *Combinatorica*, vol. 16, no. 3, pp. 417–431, 1996.
[19] M. Effros, S. E. Rouayheb, and M. Langberg, "An equivalence between network coding and index coding," *arXiv preprint arXiv:1211.6660*, 2012.
[20] S. U. Kamath, D. N. Tse, and V. Anantharam, "Generalized network sharing outer bound and the two-unicast problem," in *Network Coding (NetCod), 2011 International Symposium on.* IEEE, 2011, pp. 1–6.
[21] P. D. Seymour, "Packing directed circuits fractionally," *Combinatorica*, vol. 15, no. 2, pp. 281–288, 1995.
[22] G. Even, J. S. Naor, B. Schieber, and M. Sudan, "Approximating minimum feedback sets and multicuts in directed graphs," *Algorithmica*, vol. 20, no. 2, pp. 151–174, 1998.
[23] Z. Nutov and R. Yuster, "Packing directed cycles efficiently," in *Mathematical Foundations of Computer Science 2004.* Springer, 2004, pp. 310–321.
[24] M. A. R. Chaudhry, Z. Asad, A. Sprintson, and M. Langberg, "On the complementary index coding problem," in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on.* IEEE, 2011, pp. 244–248.
[25] S. El Rouayheb, A. Sprintson, and C. Georghiades, "On the index coding problem and its relation to network coding and matroid theory," *Information Theory, IEEE Transactions on*, vol. 56, no. 7, pp. 3187–3195, 2010.

APPENDIX

A. Proof of Theorem 3

Recall that $\mathcal{G}'$ is the directed network obtained from $\mathcal{G}$ by setting the destination node t_i to be the tail of each source link of source s_i, $i = 1, \ldots, k$. (Section III). Further, G is the (reversed) line digraph of $\mathcal{G}'$. Any set of vertices lying on a cyclic path in G corresponds to a set of edges forming a

cycle in $\mathcal{G}'$. Hence, $m - \mathsf{MAIS}(G)$ equals the cardinality of the minimum feedback edge set of the cyclic network $\mathcal{G}'$, *i.e.*, the smallest set of (unit-capacity) edges that need to be removed from $\mathcal{G}'$ to obtain an acyclic network. To show the desired result, it suffices to show any GNS cut in $\mathcal{G}$ is a feedback edge set in $\mathcal{G}'$.

Any cycle in $\mathcal{G}'$ must contain at least one of the edges connecting a destination node t_i to its source node s_i: these are the only links modified to obtain $\mathcal{G}'$ from $\mathcal{G}$, and the latter is an acyclic network. It turn, all cycles in $\mathcal{G}'$ are of the form $t_i, s_i, \ldots, t_i$.

Let S be a GNS cut in $\mathcal{G}$. By definition, there exists a permutation $\pi : [k] \to [k]$ such that if $\pi(i) \geq \pi(j)$, no path exists from s_i to t_j in $\mathcal{G} - S$. We want to show that $\mathcal{G}' - S$ is acyclic. Assume, for the sake of contradiction, that this is not the case, and let $\mathcal{C} \subseteq \{1, \ldots, k\}$ be the set of indices such that a source edge from t_i to s_i, $\forall i \in C$ lies on a cycle. Let $i^\star = \max_{i \in \mathcal{C}} \pi(i)$. Consider a cycle in $\mathcal{G}' - S$ going through an edge from $t_{i^\star}$ to $s_{i^\star}$; it must be of the form $t_{i^\star}, s_{i^\star}, \ldots, t_{i^\star}$. Without loss of generality, only one of the source edges from $t_{i^\star}$ to $s_{i^\star}$ occurs in this cycle. Since S is a GNS cut of $\mathcal{G}$, no path exists in $\mathcal{G}$ from $s_{i^\star}$ to $t_{i^\star}$. We conclude that a path from $s_{i^\star}$ to $t_{i^\star}$ in $\mathcal{G}'$ must use edges that are not available in $\mathcal{G}$, that is, edges from t_j to s_j for some $j \in \{1, \ldots, k\}$. Let $j^\star$ be the source node such that a source edge from $t_{j^\star}$ to $s_{j^\star}$ is the first source edge appearing in the path from $s_{i^\star}$ to $t_{i^\star}$. Then, the path from $s_{i^\star}$ to $t_{j^\star}$ uses only edges in $\mathcal{G} - S$ (otherwise it would go through another source edge contradicting the fact that edge from $t_{j^\star}$ to $s_{j^\star}$ is the first source edge in the cycle after $s_{i^\star}$). Hence, there is a path in $\mathcal{G} - S$ from $s_{i^\star}$ to $t_{j^\star}$, with $\pi(i^\star) > \pi(j^\star)$ which is a contradiction.

B. Proof of Theorem 4

Recall that G is the (reversed) line graph of $\mathcal{G}'$, the cyclic network obtained from $\mathcal{G}$ by connecting each destination node t_i to the source links of the source node s_i, $i = 1 \ldots, k$, as described in Section III-A.

The quantity $m - \mathsf{MAIS}(G)$ is the cardinality of the minimum feedback vertex set in G, which in turn equals the cardinality of the minimum feedback edge set (FES) in $\mathcal{G}'$, *i.e.*,

$$m - \mathsf{MAIS}(G) = \min_{\mathcal{F}' \text{ is a FES in } \mathcal{G}'} |\mathcal{F}'|. \tag{2}$$

We will show that the right hand side of (2) is equal to $\mathsf{GNSCUT}(\widetilde{\mathcal{G}})$, *i.e.*, the cardinality of the smallest GNS cut in $\widetilde{\mathcal{G}}$. In fact, we will show that for every FES $\mathcal{F}'$ in $\mathcal{G}'$, there exists a GNS cut-set $\widetilde{\mathcal{F}}$ in $\widetilde{\mathcal{G}}$ with $|\widetilde{\mathcal{F}}| = |\mathcal{F}'|$, and vice versa.

Ignoring the source links in $\widetilde{\mathcal{G}}$ (links with head vertex $\widetilde{s}_i$ for some $i \in [k]$ and no tail vertex), we consider the following trivial one-to-one mapping $\mathcal{M}$ between the links of $\widetilde{\mathcal{G}}$ and those of $\mathcal{G}'$:

- Each of the $\mathsf{MINCUT}(s_i, t_i)$ links from t_i to s_i in $\mathcal{G}'$ is mapped to a link from $\widetilde{s}_i$ to s_i in $\widetilde{\mathcal{G}}$.
- All remaining links are common in both networks.

Consider an arbitrary FES $\mathcal{F}'$ in $\mathcal{G}'$. Let $\widetilde{\mathcal{F}}$ be the image of $\mathcal{F}'$ under the mapping $\mathcal{M}$. We will show that $\widetilde{\mathcal{F}}$ is a GNS cut-set in $\widetilde{\mathcal{G}}$.

Claim 1. *Consider a k-unicast network $\mathcal{G}(\mathcal{V}, \mathcal{E})$ and a subset of links $\mathcal{F} \subset \mathcal{E}$. If $\mathcal{F}$ is not a GNS cut-set in $\mathcal{G}$, then there exists a source-destination pair s_i, t_i with a path from s_i to t_i in $\mathcal{G} - \mathcal{F}$, or a sequence of $r \geq 2$ distinct indices $i_1, \ldots, i_r \in \{1, \ldots, k\}$ such source s_{i_j} has a path to destination $t_{i_{j+1}}$, for $j = 1, \ldots, r - 1$, and s_{i_r} has a path to t_{i_1} in $\mathcal{G} - \mathcal{F}$.*

The proof of Claim 1 is deferred to the end of this section. It follows from Claim 1 that if $\widetilde{\mathcal{F}}$ is not a GNS cut-set of $\widetilde{\mathcal{G}}$, then $\mathcal{G}' - \mathcal{F}'$ contains a cycle, contradicting the fact that $\mathcal{F}'$ is a FES of $\mathcal{G}'$. We conclude that $\widetilde{\mathcal{F}}$ is a GNS cut-set in $\widetilde{\mathcal{G}}$. Note that $|\widetilde{\mathcal{F}}| = |\mathcal{F}'|$. Finally, the above implies that

$$\mathsf{GNSCUT}(\widetilde{\mathcal{G}}) \leq \min_{\mathcal{F}' \text{ is a FES in } \mathcal{G}'} |\mathcal{F}'|. \tag{3}$$

Conversely, consider an arbitrary GNS cut-set $\widetilde{\mathcal{F}}$ in $\widetilde{\mathcal{G}}$. Let $\mathcal{F}'$ be the (inverse) image of $\widetilde{\mathcal{F}}$ according to the mapping $\mathcal{M}$. We will show that $\mathcal{F}'$ is an FES in $\mathcal{G}'$.

$\widetilde{\mathcal{F}}$ is a GNS cut-set. Hence, there exists a permutation $\pi : [k] \to [k]$ such that if $\pi(i) \geq \pi(j)$, no path exists from $\widetilde{s}_i$ to t_j in $\widetilde{\mathcal{G}} - \widetilde{\mathcal{F}}, \forall i, j \in [k]$. Assume, for the sake of contradiction, that $\mathcal{F}'$ is not a FES in $\mathcal{G}'$, *i.e.*, $\mathcal{G}' - \mathcal{F}'$ contains a cycle. Any cycle in $\mathcal{G}'$ has to include a link from the destination node t_i to the source node s_i, for some $i \in \{1, \ldots, k\}$, *i.e.*, it is of the form $t_i, s_i, \ldots, t_i$ including one or several source nodes. Using an argument identical to that in the proof of Theorem 3: either *(i)* the cycle contains a path from s_i to t_i, which is a path in $\mathcal{G}$, or *(ii)* $\exists j : \pi(i) > \pi(j)$ and the cycle contains a path from s_i to t_j. This in turn implies (under the mapping $\mathcal{M}$) that $\widetilde{\mathcal{G}} - \widetilde{\mathcal{F}}$ contains either a path from $\widetilde{s}_i$ to t_i, or a path from $\widetilde{s}_i$ to t_j contradicting that $\widetilde{\mathcal{F}}$ is a GNS cut-set in $\widetilde{\mathcal{G}}$. We conclude that $\mathcal{F}'$ is a FES of $\mathcal{G}'$, while by construction $|\mathcal{F}'| = |\widetilde{\mathcal{F}}|$. Finally, the above imply that

$$\min_{\mathcal{F}' \text{ is a FES in } \mathcal{G}'} |\mathcal{F}'| \leq \mathsf{GNSCUT}(\widetilde{\mathcal{G}}). \tag{4}$$

The theorem follows from (2), (3) and (4).

1) Proof of Claim 1: We prove the contrapositive statement; if $\mathcal{G} - \mathcal{F}$ contains no source-destination pair s_i, t_i such that s_i has a path to t_i, nor a sequence of $r \geq 2$ distinct indices $i_1, \ldots, i_r \in \{1, \ldots, k\}$ with the properties described in the claim, then $\mathcal{F}$ is a GNS cut.

Consider a directed graph H on k vertices labeled $1, \ldots, k$, with vertex i corresponding to the ith source-terminal pair s_i, t_i of $\mathcal{G} - \mathcal{F}$. A directed edge (i, j) from vertex i to vertex j exists in H if and only if a path from s_i to t_j exists in $\mathcal{G} - \mathcal{F}$.

By assumption, $\mathcal{G} - \mathcal{F}$ contains no source-destination pair s_i, t_i such that s_i has a path to t_i. Hence, H contains no self-loops. Further, it is straightforward to verify that H contains a cyclic path $i_1, \ldots, i_r, i_1, r \geq 2$ if and only if the sequence of indices $i_1, \ldots, i_r \in \{1, \ldots, k\}$ satisfies the property described in the claim. By assumption, no such sequence exists in $\mathcal{G} - \mathcal{F}$ either. Hence, H is acyclic and has a topological ordering, *i.e.*, a permutation $\pi : [k] \to [k]$ of the k vertices such that for all $i, j \in [k]$, if $\pi(i) \geq \pi(j)$, then no edge exists from i to j. In turn, if $\pi(i) \geq \pi(j)$, no path exists from source s_i to

destination t_j in $\mathcal{G} - \mathcal{F}$. The existence of such a permutation implies that $\mathcal{F}$ is a GNS cut of $\mathcal{G}$ (see Def. 6). $\square$

C. Proof of Theorem 5

Consider a multiple-unicasts network $\mathcal{G}$ with k sources and m (unit-capacity) edges in the set $\mathcal{E}$. Recall that $\mathcal{G}'$ is the network formed by setting the destination t_i as the tail of all source links of s_i in $\mathcal{G}$, $\forall i \in \{1, \ldots, k\}$. We define $\mathcal{L}$ to be the set of *capacitated links* of $\mathcal{G}'$: $(a, b) \in \mathcal{L}$ if and only if there exists a link from a to b in $\mathcal{G}'$ and the capacity of (a, b), denoted by $c_{a,b}$, is equal to the number of unit links from a to b in $\mathcal{G}'$.

First, we observe that to find the smallest feedback edge set in $\mathcal{G}'$, it suffices to consider the set of capacitated links $\mathcal{L}$. Consider a set $\mathcal{F}'_E \subseteq \mathcal{E}$ be a minimal feedback edge set of $\mathcal{G}'$, *i.e.*, a minimal subset of *unit-capacity* edges whose removal from $\mathcal{G}'$ yields an acyclic network. If there exist multiple links from node u to node v in $\mathcal{G}'$, then either all or none of them is included in $\mathcal{F}'_E$. To verify that, let $e, \widehat{e}$ be unit links from u to v, and $e \in \mathcal{F}'_E$, while $\widehat{e} \notin \mathcal{F}'_E$. By construction, $\mathcal{G}' - \mathcal{F}'_E$ is an acyclic network. The minimality of $\mathcal{F}'_E$ implies that $\mathcal{G}'$ contains a cycle $u', \ldots, u \xrightarrow{e} v, \ldots, u'$ whose only edge contained in $\mathcal{F}'_E$ is e. But, $\widehat{e} \notin \mathcal{F}'_E$ and hence $u', \ldots, u \xrightarrow{\widehat{e}} v, \ldots, u'$ forms a cycle in $\mathcal{G}' - \mathcal{F}'_E$, contradicting the fact that $\mathcal{F}'_E$ is a feedback edge set.

The second key observation is that since $\mathcal{G}$ is acyclic, every cycle in $\mathcal{G}'$ must include an edge from a destination node to a source node. Equivalently, all cycles in $\mathcal{G}'$ go through the set of nodes $\mathcal{S} = \{s_1, \ldots, s_k\}$, *i.e.*, the set of k source nodes. The *minimum weight subset-feedback edge set problem* (see [22]) is the problem of finding a set of minimum weight edges that cuts all cycles passing through a *specific* set of nodes. Finding a feedback edge set in $\mathcal{G}'$ is equivalent to solving the minimum weight subset-feedback edge set problem for the set of source nodes $\mathcal{S}$ (using the capacitated links $\mathcal{L}$ with weights coinciding with the corresponding capacity). The modified sphere growing approximation algorithm of [22] with input $\mathcal{G}'$ and $\mathcal{S}$ constructs a feedback edge set of weight within a $O(\log^2 |\mathcal{S}|)$ factor from that of the minimum *fractional* weighted feedback edge set of $\mathcal{G}'$, in time polynomial in $|\mathcal{L}|$. The weight of the minimum *fractional* weighted feedback edge set coincides with the *fractional cycle packing number* of G, $r_{\text{CP}}(G)$, where G is the (reversed) line graph of $\mathcal{G}'$. In other words, the aforementioned algorithm yields a feedback edge set of G with weight at most $r_{\text{CP}}(G) \cdot O(\log^2 k)$.

It is known [24] that $m - r_{\text{CP}}(G)$ equals the broadcast rate of a vector-linear index code C for index coding instance with side-information graph G. In other words, there exists an index code C, which can be defined over any field $\mathbb{F}$, that achieves broadcast rate $\beta_{\text{VL}}^{\mathbb{F}}(\mathcal{C}; G) = m - r_{\text{CP}}(G)$. The basis for the coding scheme is that every cycle in G saves a transmission. This implies $m - \beta_{\text{VL}}(G) \geq r_{\text{CP}}(G)$. Further, by duality in Theorem 1, there is a feasible code for the relaxed-correlated sources problem on G whose joint entropy rate is $r_{\text{CP}}(G)$.

D. Proof of Theorem 6

Our proof relies on the examples previously used to separate scalar linear and non-linear broadcast rates for index coding. We first show that the same examples imply separation results between vector-linear broadcast rates over two different fields. These results require the application of an *uncertainty principle* (Lemma 6) for vector-linear broadcast rates for index coding.[3] At first sight, those examples imply separation for only symmetric rates. However, we exploit certain symmetry properties that imply separation for sum-rates (asymmetric rates) when written as network coding instances. We show that that there exist multiple-unicasts network instances for which optimal sum-rate (for independent sources) over two different prime fields $\mathbb{F}_p$ and $\mathbb{F}_q$ can be separated by a factor of $k^{1-o(1)}$. This, along with Theorem 5 concludes the proof.

Index coding – optimum vector-linear broadcast rates over two different prime fields: The maximum vector-linear broadcast rate $\beta_{\text{VL}}^{\mathbb{F}}(G)$ over a field $\mathbb{F}$ for an index coding instance with side-information graph G may vary significantly with the choice of $\mathbb{F}$. We show that for any two prime fields $\mathbb{F}_p$ and $\mathbb{F}_q$, there exist index coding instances G for which $\beta_{\text{VL}}^{\mathbb{F}_p}(G)$ and $\beta_{\text{VL}}^{\mathbb{F}_q}(G)$ for which $\beta_{\text{VL}}^{\mathbb{F}_q}(G)/\beta_{\text{VL}}^{\mathbb{F}_p}(G) \geq k^{1-o(1)}$.

For any two prime fields $\mathbb{F}_p$ and $\mathbb{F}_q$, Lubetzky and Stav [12] construct an undirected graph G_u such that $\text{minrk}_{\mathbb{F}_q}(G_u) \leq k^{o(1)}$ and $\text{minrk}_{\mathbb{F}_p}(\overline{G}_u) \leq k^{o(1)}$, where $\overline{G}_u$ denotes the undirected graph complement. The set of vertices of G_u is the collection of all s-subsets of $[r]$. An undirected edge between two vertices X and Y exists if and only if $X \neq Y$ and $|X \bigcap Y| \equiv -1(\mod p^b)$. The parameters r, s and b depend on the choice of p and q (see Section 2 in [12]). Recall that $\text{minrk}_{\mathbb{F}}(G_u)$ equals the optimal broadcast rate on the index coding instance with side-information graph G_u among *scalar* linear codes over $\mathbb{F}$. We conclude that

$$\beta_{\text{VL}}^{\mathbb{F}_p}(G_u) \leq k^{o(1)} \quad \text{and} \quad \beta_{\text{VL}}^{\mathbb{F}_q}(\overline{G}_u) \leq k^{o(1)}. \qquad (5)$$

By Lemma 6 (*uncertainty principle*), we know that for any field $\mathbb{F}_p$, $\beta_{\text{VL}}^{\mathbb{F}_p}(\overline{G}_u) \cdot \beta_{\text{VL}}^{\mathbb{F}_p}(G_u) \geq k$. This implies,

$$\beta_{\text{VL}}^{\mathbb{F}_q}(\overline{G}_u) \leq k^{o(1)} \quad \text{and} \quad \beta_{\text{VL}}^{\mathbb{F}_p}(\overline{G}_u) \geq k^{1-o(1)}. \qquad (6)$$

The network coding instance: Given $\overline{G}_u$, we consider the multiple-unicasts network $\overline{\mathcal{G}}_u$ with k sources constructed as follows[4]:

1) Create a source-destination pair s_i, t_i for each vertex i in $\overline{G}_u$. A unit-capacity link from s_i to t_j exists in the network if and only if a (directed) edge (i, j) exists in $\overline{G}_u$. (Undirected edges in $\overline{G}_u$ are equivalent to two directed ones.)

2) Introduce two nodes a and b with a unit-capacity link from a to b. Connect s_i to a and b to t_i with unit-capacity links, $\forall i \in \{1, \ldots, k\}$.

---

[3]We note that Lemma 6, although implied by results in [11], is to the best of our knowledge not explicitly stated in the literature.

[4]This is a slight modification of the construction in (Fig. 2)[25].

3) Introduce a single unit-capacity source link (with no tail) for each s_i. Note that $\mathsf{MINCUT}(s_i, t_i) = 1$ since there is a single path from s_i to t_i; through link (a, b).

Any index code achieving broadcast rate β in the index coding instance with side-information graph $\overline{G}_u$, corresponds to a symmetric rate point $(1/\beta, 1/\beta, \ldots)$ on the multiple-unicasts network instance $\overline{\mathcal{G}}_u$ with independent sources, and vice versa [25].

In principle, the optimum sum-rate on the multiple-unicasts network $\overline{\mathcal{G}}_u$ could be achievable by an asymmetric rate point. We show that for $\overline{\mathcal{G}}_u$ this is not the case; the optimal sum-rate can be achieved by a symmetric rate point. Assume that an asymmetric rate point $(r_1, \ldots, r_k)$ is achievable on $\overline{\mathcal{G}}_u$. The original graph $\overline{G}_u$ is vertex transitive: given any two vertices v_1 and v_2 of $\overline{G}_u$, there is some automorphism $f : V(G) \to V(G)$ such that $f(v_1) = v_2$. The multiple-unicasts network corresponding to the image of $\overline{G}_u$ under $f(\cdot)$ is identical to $\overline{\mathcal{G}}_u$ (with rearranged sources), implying that the rate point $\left(r_{f^{-1}(1)}, \ldots, r_{f^{-1}(k)}\right)$ is also achievable on $\overline{\mathcal{G}}_u$. $\overline{G}_u$ has a special structure: Each vertex is an s-subset of $[r]$ and a vertex for each subset exists. Further, the edge between two vertices is determined by only the size of the intersection of subsets corresponding to the pair of vertices. Therefore, the number of automorphisms f mapping v to u in $\overline{G}_u$, is equal to the number of automorphisms mapping v' to u in $\overline{G}_u$, and so on. This is because, all subsets in some sense are equivalent and therefore can be mapped in similar ways to a given subset. Consider the entire collection of achievable rate tuples corresponding to all the automorphisms. Then, the number of rate tuples in which user u gets rate r_v is equal to the number of tuples in which he gets rate r'_v and so on. The same holds for all users. Therefore, time sharing among those rate tuples yields a symmetric rate point achieving sum-rate equal to $r_1 + \ldots + r_k$.

In summary, any index code achieving broadcast rate β in the index coding instance with side-information graph $\overline{G}_u$, corresponds to a symmetric rate point in the multiple-unicasts network coding instance on $\overline{\mathcal{G}}_u$ with independent sources. Further, the optimum sum-rate on $\overline{\mathcal{G}}_u$ is achievable by a symmetric rate point.

We conclude that there exists a symmetric rate tuple for $\overline{\mathcal{G}}_u$ that achieves sum-rate $R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_p) = k/\beta_{\mathrm{VL}}^{\mathbb{F}_p}(\overline{G}_u)$ and a tuple achieving $R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_q) = k/\beta_{\mathrm{VL}}^{\mathbb{F}_q}(\overline{G}_u)$. In conjunction with (6), we conclude that

$$R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_p) \leq k^{o(1)} \text{ and } R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_q) \geq k^{1-o(1)}, \quad (7)$$

which proves the second part of the theorem.

Contradiction: Assume for the sake of contradiction that $R^{\mathrm{CO}}(\overline{\mathcal{G}}_u, \mathbb{F}_p) < k^{1-\delta} \cdot R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_p)$. By Theorem 5, we know that for any field $\mathbb{F}$, $r_{\mathrm{CP}}(\overline{G}_u) \leq R^{\mathrm{CO}}(\overline{\mathcal{G}}_u, \mathbb{F}) \leq r_{\mathrm{CP}}(\overline{G}_u) \cdot O(\log^2 k)$. Then, according to our assumption,

$$r_{\mathrm{CP}}(\overline{G}_u) < k^{1-\delta} \cdot R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_p). \quad (8)$$

Also, $R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_q) \leq R^{\mathrm{CO}}(\overline{\mathcal{G}}_u, \mathbb{F}_q)$ which implies that

$$R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_q) \leq r_{\mathrm{CP}}(\overline{G}_u) \cdot O(\log^2 k). \quad (9)$$

Combinining (8) and (9), we obtain

$$\begin{aligned} R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_q) &< k^{1-\delta} \cdot O(\log^2 k) \cdot R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_p) \\ &< k^{1-o(1)} \cdot R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_p), \end{aligned}$$

which contradicts (7).

Conclusion: For any prime field $\mathbb{F}_p$, we can construct a network $\overline{\mathcal{G}}_u$ such that $R^{\mathrm{CO}}(\overline{\mathcal{G}}_u, \mathbb{F}_p) \geq k^{1-\delta} \cdot R^{\mathrm{MU}}(\overline{\mathcal{G}}_u, \mathbb{F}_p)$. This proves the first part of the theorem.

E. Lower bounds on vector-linear Index Coding using strong products of graphs

Definition 7. (*Strong product of digraphs*) *The strong product $G \otimes H$ of two digraphs G and H is a digraph on the set of vertices $V(G \otimes H) \triangleq \{(u, v) : u \in V(G), v \in V(H)\}$, with set of edges $E(G \otimes H)$ that contains an edge from (u, v) to (u', v') if and only if the following conditions both hold:*
1) $u = u' \vee (u, u') \in E(G)$
2) $v = v' \vee (v, v') \in E(H)$.

Definition 8. (*Complement of a digraph*) *The complement of a digraph $G = (V, E)$ is a digraph $\overline{G} = (V, \overline{E})$ on V, where $\overline{E} = \{(u, v) : u, v \in V, (u, v) \notin E\}$.*

$\alpha(G)$ denotes an independent set of G, i.e. a set of vertices any two of which have no edge in either direction.

Definition 9. (*k-blowup of a digraph*) *The k-blowup of a digraph G on n vertices, denoted by $G[k]$, is a digraph on $k \cdot n$ vertices such that:*
- *$V(G[k])$ contains k distinct vertices $v^{(1)}, \ldots, v^{(k)}$ for each vertex $v \in V(G)$.*
- *An edge from $u^{(i)}$ to $v^{(j)}$, $i, j \in [k]$ exists in $E(G[k])$ if and only if $(u, v) \in E(G)$.*

The definition of the k-blowup implies a surjective mapping from $V(G[k])$ to $V(G)$. Throughout this paper, we use the convention that vertices $v^{(i)}$, $i = 1, \ldots, k$, in $G[k]$ originate from (map to) vertex v in G. Conversely, v yields vertices $v^{(1)}, \ldots, v^{(k)}$ in $G[k]$. Note that $\{v^{(i)}\}_{i=1}^k$ forms an independent set of cardinality k in $G[k]$. Further, each edge $(u, v) \in E(G)$ corresponds to a directed biclique[5] $(\{u^{(i)}\}_{i=1}^k, \{v^{(i)}\}_{i=1}^k)$ in $G[k]$.

Definition 10. $\mathsf{MAIS}(G)$ *of a directed graph G is the cardinality of the largest set $\widehat{V} \subseteq V(G)$, such that the subgraph of G induced by $\widehat{V}$ is acyclic.*

Remark 4. *For any undirected graph G, $\mathsf{MAIS}(G) = \alpha(G)$. An undirected graph G can be considered as a digraph such that $(u, v) \in E(G) \Rightarrow (v, u) \in E(G)$. Hence, any edge (u, v) in G corresponds to a cycle $u \to v \to u$. In undirected graphs, acyclic subgaphs coincide with independent sets.*

Let G be a digraph on n vertices without self-loops. We say that an $n \times n$ matrix $\mathbf{A}$ over a finite field $\mathbb{F}_q$ *fits* G if (i) $A_{ii} \in \mathbb{F}_q \backslash \{0\}$, $\forall i \in [n]$, and (i) $A_{ij} = 0$ $\forall (i, j) \notin E(G)$. Let $\mathrm{rank}_q(\mathbf{A})$ denote the rank of $\mathbf{A}$ over $\mathbb{F}_q$.

---
[5] A directed biclique is a biclique with partition (U, V) whose edges have been oriented from U to V.

Definition 11. $\mathsf{minrk}_q(G) = \min\{\mathrm{rank}_q(\mathbf{A}) \,:\, \mathbf{A} \text{ fits } G\}$.

Lemma 2. *Consider m directed graphs G_i, $i = 1, \ldots, m$, and m positive integers $k_i \in \mathbb{N}_+$. Let G_α be the strong product of $G_i[k_i]$, $i = 1, \ldots, m$, i.e.,*

$$G_\alpha = \otimes_{i=1}^m G_i[k_i],$$

and G_β be the $(\prod_{i=1}^m k_i)$-blowup of the strong product of G_i's:

$$G_\beta = (\otimes_{i=1}^m G_i)[\textstyle\prod_{i=1}^m k_i].$$

Then, G_α is (isomorphic to) a spanning subgraph of G_β.

Proof: We establish a bijection $f : V(G_\alpha) \to V(G_\beta)$ between the vertices of G_α and G_β, and show that two vertices $u, v \in V(G_\alpha)$ are connected with an edge only if vertices $f(u), f(v)$ in G_β are also connected.

Let n_i denote the number of vertices in G_i, and $N \triangleq \prod_{i=1}^m n_i$. By construction, $G_i[k_i]$ is a graph on $k_i \cdot n_i$ vertices: k_i vertices $v_i^{(j)}$, $j \in \{1, \ldots, k_i\}$, for each $v_i \in V(G_i)$. The strong product $G_\alpha = \otimes_{i=1}^m G_i[k_i]$ is a graph on $K \cdot N$ vertices, where $K \triangleq \prod_{i=1}^m k_i$. The vertex set is formed as the cartesian product $V(G_1[k_1]) \times \cdots \times V(G_m[k_m])$: the set of all m-tuples $(v_1^{(j_1)}, \ldots, v_m^{(j_m)})$ whose ith entry is a vertex of $G_i[k_i]$.

Similarly, the strong product $\otimes_{i=1}^m G_i$ is a graph on a set of N vertices. Each vertex corresponds to an m-tuple $(v_1, \ldots, v_m)$, where $v_i \in V(G_i)$. Its K-blowup is a graph G_β on $K \cdot N$ vertices: K vertices $(v_1, \ldots, v_m)^{(j_1, \ldots, j_m)}$, $j_i \in \{1, \ldots, k_i\}$ for each vertex $(v_1, \ldots, v_m) \in V(\otimes_{i=1}^m G_i)$. Observe that we use an m-tuple $(j_1, \ldots, j_m)$ to enumerate the K vertices in G_β originating from a single vertex of the strong product $\otimes_{i=1}^m G_i$. This alternative labeling establishes the one-to-one mapping between the vertices of G_α and G_β:

$$f\big((v_1^{(j_1)}, \ldots, v_m^{(j_m)})\big) = (v_1, \ldots, v_m)^{(j_1, \ldots, j_m)}.$$

It remains to show that two vertices $u = (u_1^{(j_1)}, \ldots, u_m^{(j_m)})$ and $v = (v_1^{(l_1)}, \ldots, v_m^{(l_m)})$ in G_α are connected with an edge only if $f(u), f(v)$ in G_β are also connected. By construction, an edge between u and v exists if an only if $\forall\, i \in \{1, \ldots m\}$:

$$u_i^{(j_i)} = v_i^{(l_i)} \quad \vee \quad \big(v_i^{(j_i)}, v_i^{(l_i)}\big) \in E(G_i[k_i]). \quad (10)$$

Note that, $u_i^{(j_i)} = v_i^{(l_i)} \Rightarrow u_i = v_i \in V(G_i)$. Further, an edge between $u_i^{(j_i)}$ and $v_i^{(l_i)}$ exists in $G_i[k_i]$ if and only if an edge between u_i and v_i exists in G_i. That is,

$$\big(u_i^{(j_i)}, v_i^{(l_i)}\big) \in E\left(G_i[k_i]\right) \quad \Leftrightarrow \quad (u_i, v_i) \in E(G_i).$$

Hence, the existence of an edge between u and v in G_α, implies that $\forall i \in \{1, \ldots m\}$,

$$u_i = v_i \ \vee \ (u_i, v_i) \in E(G_i), \quad u_i, v_i \in V(G_i).$$

It follows that $\otimes_{i=1}^m G_i$ contains an edge between $(u_1, \ldots, u_m)$ and $(v_1, \ldots, v_m)$, and in turn G_β contains an edge between $f(u) = (u_1, \ldots, u_m)^{(j_1, \ldots, j_m)}$ and $f(v) = (v_1, \ldots, v_m)^{(l_1, \ldots, l_m)}$, which completes the proof. ∎

Corollary 2. *For the graphs G_α and G_β of Lemma 2,*

1) $\alpha(G_\alpha) \geq \alpha(G_\beta)$,
2) $\mathsf{MAIS}(G_\alpha) \geq \mathsf{MAIS}(G_\beta)$.

Proof: The corollary is a straightforward consequence of the fact that G_α is a spanning subgraph of G_β:

Part (i): Consider the maximum independent set $\mathcal{I}(G_\beta)$ of G_β. Since G_α is a spanning subgraph of G_β, no edge exists between the vertices in the inverse image $f^{-1}(\mathcal{I}(G_\beta))$ of $\mathcal{I}(G_\beta)$. In other words, $f^{-1}(\mathcal{I}(G_\beta))$ is an independent set in G_α. Given that $f(\cdot)$ is a bijection, we conclude that $|\mathcal{I}(G_\alpha)| \geq |f^{-1}(\mathcal{I}(G_\beta))| = \mathcal{I}(G_\beta)$.

Part (ii): Let $\mathcal{A}(G)$ denote the largest subset of vertices in G such that the subgraph induced by $\mathcal{A}(G)$ is acyclic. Consider the vertices of G_α in the inverse image of $\mathcal{A}(G_\beta)$, $f^{-1}(\mathcal{A}(G_\beta))$. The corresponding subgraph of G_α contains no cycle. Given that $f(\cdot)$ is a bijection, we conclude that $\mathsf{MAIS}(G_\alpha) \geq |f^{-1}(\mathcal{A}(G_\beta))| = \mathcal{A}(G_\beta)$. ∎

Lemma 3. *For any two directed graphs G and H,*

$$\mathsf{MAIS}(G \otimes H) \geq \mathsf{MAIS}(G) \cdot \mathsf{MAIS}(H).$$

Proof: Let $\mathcal{A}(G) \subseteq V(G)$ denote the largest set of vertices such that the subgraph of G induced by $\mathcal{A}(G)$ is acyclic, i.e., $|\mathcal{A}(G)| = \mathsf{MAIS}(G)$. Consider the set

$$\mathcal{S} \triangleq \{(u, v) : u \in \mathcal{A}(G), v \in \mathcal{A}(H)\} \subseteq V(G \otimes H).$$

We show that the subgraph of $G \otimes H$ induced by $\mathcal{S}$ is acyclic. Assume for the sake of contradiction that there exists a subset $\mathcal{C} \subseteq \mathcal{S}$ of $m > 1$ vertices that lie on a directed cycle: $(u_0, v_0) \to \cdots \to (u_{m-1}, v_{m-1}) \to (u_0, v_0)$. The existence of edges between consecutive vertices implies that the following is true for all $i \in \{0, \ldots, m-1\}$:

$$u_i = u_{i+1 \bmod m} \ \vee \ (u_i, \, u_{i+1 \bmod m}) \in E(G).$$

Hence, the sequence of vertices $u_0, \ldots, u_{m-1}, u_0$ in $\mathcal{A}(G)$ either consists of a single vertex repeating $m + 1$ times, or forms a cyclic path in G, ignoring transitions where consecutive vertices in the sequence are the same vertex in G. The same holds for the sequence of vertices $v_0, \ldots, v_{m-1}, v_0$ in $\mathcal{A}(H)$. At least one of the two sequences must contain more than one distinct vertices; otherwise, $|\mathcal{C}| = 1$. Therefore, either the subgraph of G induced by $\mathcal{A}(G)$ or the subgraph of H induced by $\mathcal{A}(H)$ contains a cycle, contradicting the definition of $\mathcal{A}(\cdot)$. We conclude that the subgraph of $G \otimes H$ induced by $\mathcal{S}$ is acyclic and $\mathsf{MAIS}(G \otimes H) \geq |\mathcal{S}| = \mathsf{MAIS}(G) \cdot \mathsf{MAIS}(H)$. ∎

Lemma 4. *For any directed graph G and its k-blowup $G[k]$,*

1) $\alpha(G[k]) = k \cdot \alpha(G)$,
2) $\mathsf{MAIS}(G[k]) = k \cdot \mathsf{MAIS}(G)$.

Proof: By construction $G[k]$ contains k distinct unconnected vertices $v^{(1)}, \ldots, v^{(k)}$ for each vertex $v \in V(G)$.

Part (i): Consider the set

$$\mathcal{S} \triangleq \left\{ v \in V(G) : v^{(i)} \in \mathcal{I}(G[k]), i \in [k] \right\}. \qquad (11)$$

$\mathcal{S}$ is an independent set in G: if an edge exists between two vertices $u, v \in \mathcal{S}$ in G, then $(v^{(i)}, u^{(j)}) \in E(G[k])$, $\forall i, j \in [k]$, contradicting the fact that $v^{(i)}, u^{(j)} \in \mathcal{I}(G[k])$ for some $i, j \in [k]$. Taking into account that at most k vertices in $\mathcal{I}(G[k])$ (and generally in $G[k]$) originate from a single vertex in G, we conclude that

$$\alpha(G) \geq |\mathcal{S}| = \frac{1}{k} \cdot \alpha(G[k]). \qquad (12)$$

$\mathcal{T} \triangleq \left\{ v^{(i)} \in V(G[k]), \ i = 1, \ldots, k : v \in \mathcal{I}(G) \right\}$ is an independent set in $G[k]$. By construction, no edge exists between $v^{(i)}$ and $v^{(j)}$, $i \neq j$. Further, an edge between two vertices $v^{(i)}, u^{(j)} \in \mathcal{T}$ in $G[k]$ exists only if an edge exists between v and u in G. But $v, u \in \mathcal{I}(G)$. Thus,

$$\alpha(G[k]) \geq |\mathcal{T}| = k \cdot \alpha(G). \qquad (13)$$

The desired result follows from (12) and (13).

Part (ii): Let $\mathcal{A}(G) \subseteq V(G)$ denote the largest set of vertices such that the subgraph of G induced by $\mathcal{A}(G)$ is acyclic, *i.e.*, $\mathsf{MAIS}(G) = |\mathcal{A}(G)|$. Consider the set

$$\mathcal{S} \triangleq \left\{ v \in V(G) : v^{(i)} \in \mathcal{A}(G[k]), i \in [k] \right\}.$$

The subgraph of G induced by $\mathcal{S}$ contains no cycle. To verify that, observe that if there exists a subset $\mathcal{C} \subseteq \mathcal{S}$ of m vertices forming a directed cycle $v_1 \to \cdots \to v_m \to v_1$, then $v_1^{(i_1)} \to \cdots \to v_m^{(i_m)} \to v_1^{(i_1)}$ forms a cycle in $G[k]$ for any combination of $i_1, \ldots, i_m$. Therefore, a cycle in $G[k]$ can be formed using vertices in $\mathcal{A}(G[k])$, contradicting the definition of $\mathcal{A}(\cdot)$. It follows that

$$\mathsf{MAIS}(G) \geq |\mathcal{S}| = \frac{1}{k} \cdot \mathsf{MAIS}(G[k]). \qquad (14)$$

Conversely, the set

$$\mathcal{T} \triangleq \left\{ v^{(i)} \in V(G[k]), \ i = 1, \ldots, k : v \in \mathcal{A}(G) \right\}$$

induces an acyclic subgraph in $G[k]$: if there exist a subset of vertices in $\mathcal{T}$ forming a cycle $v^{(i)} \to u^{(j)} \to \cdots \to v^{(i)}$ in $G[k]$, then $v \to u \to \cdots \to v$ is a cycle in G, contradicting the fact that $v, u, \ldots \in \mathcal{A}(G)$. Thus,

$$\mathsf{MAIS}(G[k]) \geq |\mathcal{T}| = k \cdot \mathsf{MAIS}(G). \qquad (15)$$

The desired result follows from (14) and (15). $\blacksquare$

Lemma 5. *For any $k, m \in \mathbb{N}_+$,*

$$\frac{1}{k} \cdot \mathsf{minrk}_q (G[k]) \geq \mathsf{MAIS}(G^m)^{1/m}.$$

Proof: We have

$$\mathsf{minrk}_q^m(G[k]) \overset{(\alpha)}{\geq} \mathsf{minrk}_q(\otimes^m(G[k])) \overset{(\beta)}{\geq} \mathsf{MAIS}(\otimes^m(G[k]))$$
$$\overset{(\gamma)}{\geq} \mathsf{MAIS}((\otimes^m G)[k^m]) \overset{(\delta)}{\geq} k^m \cdot \mathsf{MAIS}(\otimes^m G),$$

where (α) follows from the submultiplicativity of the $\mathsf{minrk}_q(\cdot)$ function [12], (β) from [14], (γ) follows from Corollary 2, and (δ) from Lemma 4. $\blacksquare$

Lemma 6. $\beta_{\mathrm{VL}}^{\mathbb{F}}(G) \cdot \beta_{\mathrm{VL}}^{\mathbb{F}}(\overline{G}) \geq n.$

Proof: For any $k_1, k_2 \in \mathbb{N}_+$,

$$\mathsf{minrk}_{\mathbb{F}}(G[k_1]) \cdot \mathsf{minrk}_{\mathbb{F}}(\overline{G}[k_2])$$
$$\overset{(\alpha)}{\geq} \mathsf{minrk}_{\mathbb{F}}(G[k_1] \otimes \overline{G}[k_2]) \overset{(\beta)}{\geq} \alpha(G[k_1] \otimes \overline{G}[k_2])$$
$$\overset{(\gamma)}{\geq} \alpha((G \otimes \overline{G})[k_1 \cdot k_2]) \overset{(\delta)}{\geq} k_1 \cdot k_2 \cdot \alpha(G \otimes \overline{G})$$
$$\geq k_1 \cdot k_2 \cdot n, \qquad (16)$$

where (α) follows from the submultiplicativity of the minrk_q function [12], (β) from Lemma 5 and the fact that $\mathsf{MAIS}(G) \geq \alpha(G)$, (γ) from Corollary 2, and (δ) from Lemma 4. The last inequality follows from the fact that $\alpha(G \otimes \overline{G}) \geq n$. To verify that, observe that the set $\{(u, u) : u \in V(G)\} \subseteq V(G \times \overline{G})$ forms an independent set of cardinality n in $G \times \overline{G}$. Inequality (16) holds for all k_1 and k_2. The desired result follows from Theorem 7. $\blacksquare$

Extending the arguments for optimality of scalar minrk in [14] for scalar linear index coding, it can be shown that the optimum broadcast rate of vector-linear index codes over all fields can be shown to be infimum of the normalized minranks of graph blowups.

Theorem 7 ([11]). $\beta_{\mathrm{VL}}^{\mathbb{F}}(G) = \inf_k \frac{\mathsf{minrk}_{\mathbb{F}}(G[k])}{k}.$

Theorem 8. [6] $\Theta(G) \leq \beta_{\mathrm{VL}}(G).$

Proof: For any $k, m \in \mathbb{N}_+$,

$$\alpha(\otimes^m G[k]) \overset{\mathrm{Cor.}\ 2}{\geq} \alpha((\otimes^m G)[k^m]) \overset{\mathrm{Lem.}\ 4}{=} k^m \alpha(\otimes^m G).$$

Hence, by the definition of Shannon capacity,

$$\Theta(G) = \lim_{m \to \infty} \alpha(\otimes^m G)^{1/m} \leq \lim_{m \to \infty} \frac{1}{k} \alpha(\otimes^m G[k])^{1/m}$$
$$= \frac{1}{k} \Theta(G[k]) \leq \frac{1}{k} \mathsf{minrk}_{\mathbb{F}}(G[k]), \qquad \forall k \in \mathbb{N}_+.$$

The desired result follows from Theorem 7. $\blacksquare$

Theorem 9. *For any graph G and $m \in \mathbb{N}_+$,*
1) $\alpha(\otimes^m G)^{1/m} \leq \Theta(G) \leq \beta_{\mathrm{VL}}(G).$
2) $\mathsf{MAIS}(\otimes^m G)^{1/m} \leq \beta_{\mathrm{VL}}(G).$

Proof: For the second part, by Theorem 7, and Lemma 5, $\beta_{\mathrm{VL}}^{\mathbb{F}} \geq \mathsf{MAIS}(\otimes^m G)^{1/m}$, for any field $\mathbb{F}$. The desired result follows taking infimum over all fields. The first part follows from the fact that $\alpha(\otimes^m G) \leq \mathsf{MAIS}(\otimes^m G)$, in conjunction with the definition of $\Theta(G)$ and Theorem 8. $\blacksquare$

---

[6][11] contains proofs for Theorems 8 and 9(i). Lemma 6 can be also inferred, but is not explicitly stated. We give a different proof based on strong products of graphs that also generalizes bounds by tensorizing MAIS(G) instead of just tensorizing $\alpha(G)$.